Smudge Attacks on Smartphone Touch Screens

Adam J. Aviv, Katherine Gibson, Evan Mossop, Matt Blaze, and Jonathan M. Smith
Department of Computer and Information Science – University of Pennsylvania
{aviv,gibsonk,emossop,blaze,jms}@cis.upenn.edu

Abstract

Touch screens are an increasingly common feature on personal computing devices, especially smartphones, where size and user interface advantages accrue from consolidating multiple hardware components (keyboard, number pad, *etc.*) into a single software definable user interface. Oily residues, or *smudges*, on the touch screen surface, are one side effect of touches from which frequently used patterns such as a graphical password might be inferred.

In this paper we examine the feasibility of such *smudge attacks* on touch screens for smartphones, and focus our analysis on the Android password pattern. We first investigate the conditions (*e.g.*, lighting and camera orientation) under which smudges are easily extracted. In the vast majority of settings, partial or complete patterns are easily retrieved. We also emulate usage situations that interfere with pattern identification, and show that pattern smudges continue to be recognizable. Finally, we provide a preliminary analysis of applying the information learned in a smudge attack to guessing an Android password pattern.

1 Introduction

Personal computing devices now commonly use touch screen inputs with application-defined interactions that provide a more intuitive experience than hardware keyboards or number pads. Touch screens are touched, so oily residues, or *smudges*, remain on the screen as a side effect. Latent smudges may be usable to infer recently and frequently touched areas of the screen – a form of information leakage.

This paper explores the feasibility of *smudge attacks*, where an attacker, by inspection of smudges, attempts to extract sensitive information about recent user input. We provide initial analysis of the capabilities of an attacker who wishes to execute a smudge attack. While this analysis is restricted to smartphone touch screens, specifically attacks against the Android password pattern, smudge attacks may apply to a significantly larger set of devices, ranging from touch screen ATMs and DRE voting machines to touch screen PIN entry systems in convenience stores.

We believe smudge attacks are a threat for three reasons. First, smudges are surprisingly[1] persistent in time. Second, it is surprisingly difficult to incidentally obscure or delete smudges through wiping or pocketing the device. Third and finally, collecting and analyzing oily residue smudges can be done with readily-available equipment such as a camera and a computer[2].

To explore the feasibility of smudge attacks against the Android password pattern, our analysis begins by evaluating the conditions by which smudges can be photographically extracted from smartphone touch screen surfaces. We consider a variety of lighting angles and light sources as well as various camera angles with respect to the orientation of the phone. Our results are extremely encouraging: in one experiment, the pattern was partially identifiable in 92% and fully in 68% of the tested lighting and camera setups. Even in our worst performing experiment, under less than ideal pattern entry conditions, the pattern can be partially extracted in 37% of the setups and fully in 14% of them.

We also consider simulated user usage scenarios based on expected applications, such as making a phone call, and if the pattern entry occurred prior to or post application usage. Even still, partial or complete patterns are easily extracted. We also consider incidental contact with clothing, such as the phone being placed in a pocket; information about the pattern can still be retrieved. Finally, we provide preliminary analysis of applying a smudge attack to the Android password pattern and how the information learned can be used to guess likely passwords.

Next, in Sec. 2, we provide our threat model, followed by background on the Android password pattern in Sec. 3. Our experimental setup is presented in Sec. 4, including a primer on lighting and photography. Experimental results are presented in Sec. 5, and a discussion of applying a smudge attack to the Android pattern password is presented in Sec. 6. Related work is provided in Sec. 7, and we conclude in Sec. 8.

[1] One smartphone in our study retained a smudge for longer than a month without any significant deterioration in an attacker's collection capabilities.

[2] We used a commercial photo editing package to adjust lighting and color contrast, only, but software included with most operating systems is more than sufficient for this purpose.

2 Threat Model

We consider two styles of attacker, passive and active. A *passive attacker* operates at a distance, while an *active attacker* has physical control of the device.

A passive attacker who wishes to collect smartphone touch screen smudges may control the camera angle, given the attacker controls the camera setup, but the smartphone is in possession of its user. The attacker has no control of the places the user takes the smartphone, and thus cannot control lighting conditions or the angle of the phone with respect to the camera. The attacker can only hope for an opportunity to arise where the conditions are right for good collection. An active attacker, however, is capable of controlling the lighting conditions and is allowed to alter the touch screen to increase retrieval rate. This could include, for example, cleaning the screen prior to the user input, or simply moving the touch screen to be at a particular angle with respect to the camera.

For the purposes of our experiment, we make a strong assumption about the attacker's "activeness;" she is in possession of the device, either surreptitiously or by confiscation, and is capable of fully controlling the lighting and camera conditions to extract information. We believe such an attacker is within reason considering search and seizure procedures in many countries and states. However, a passive smudge attack, *e.g.*, via telephotography, can still be useful in a later active attack, where the touch screen device becomes available. The information obtained will often still be fresh – users tend to leave their passwords unchanged unless they suspect a compromise [3] – encouraging multiphase attack strategies.

3 Android Password Pattern

The Android password pattern is one of two unlock mechanisms, as of the release of Android 2.2 where alpha-numeric pins are now allowed [1]. However, the password pattern is currently the primary authentication mechanism on the vast majority of Android devices that have not yet received the update, and the pattern remains an authentication option on Android 2.2 devices.

The Android pattern is one style of graphical passwords where a user traverses an onscreen 3x3 grid of contacts points. A pattern can take on a number of shapes and can be defined as an ordered list of contact points (Fig. 1 provides an indexing scheme). For example, the "L" shaped password can be represented as the ordered list |14789|, *i.e.*, the user begins by touching contact point 1, drawing downward towards point 7, and finally across to point 9[3].

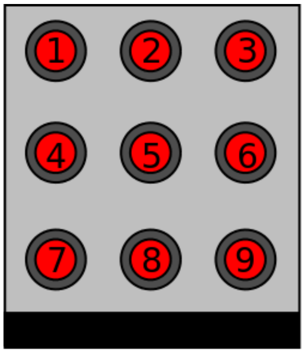

Figure 1: An illustration of the Android password pattern screen with overlaid identification numbers on contact points.

There are a three restrictions on acceptable patterns. It must contact a minimum of four points, so a single stroke is unacceptable. Additionally, a contact point can only be used once. These two restrictions imply that every pattern will have at least one direction change, and as the number of contact points increases, more and more such direction changes are required. Such convoluted connections of smudges may actually increase the contrast with background noise, as one of our experiments suggests (see Sec. 5).

The last, and most interesting, restriction applies to intermediate contact points: If there exists an intermediate point between two other contact points, it must also be a contact point, unless, that point was previously contacted. For example, in the "L" shaped pattern, it must always contain points 4 and 8 even though the ordered list |179| would construct the exact same pattern If a user attempted to avoid touching either point 4 or 8, both would be automatically selected. Conversely, consider a "+" shaped pattern constructed by either the order list |25846| or |45628|, the connected points |46| or |28| are allowed because point 5 was previously contacted.

Due to the intermediate contact point restriction, the password space of the Android password pattern contains 389,112 possible patterns[4]. This is significantly smaller than a general ordering of contact points, which contains nearly 1 million possible patterns. Still, this is a reasonably large space of patterns, but when considering information leakage of smudge attacks, an attacker can select a highly likely set of patterns, increasing her chances of guessing the correct one before the phone locks-out[5]. Sometimes, even, the precise pattern can be determined.

[3] Although a pattern can be entered using two fingers, stepping in order to simulate a drag from dot-to-dot, it is unlikely common practice because it requires more effort on the part of the user and is not part of the on-screen instructions provided by Android.

[4] Due to the complexity of the intermediate contact point restriction, we calculated this result via brute force methods.

[5] Android smartphones require the user to enter a Google user-name and password to authenticate after 20 failed pattern entry attempts.

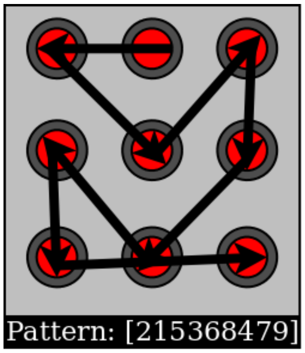

Figure 2: Password pattern used for captures; it contains streaks in all orientations and most directions.

4 Experimental Setup

In this section we present our experimental setup for capturing smudges from smartphone touch screens, including a background on photography and lighting. We experimented with two Android smartphones, the HTC G1 and the HTC Nexus1, under a variety of lighting and camera conditions. We also experimented with simulated phone application usage and smudge distortions caused by incidental clothing contact.

4.1 Photography and Lighting

This paper primarily investigates the camera angles and lighting conditions under which latent "smudge patterns" can be recovered from touchscreen devices. The fundamental principles of lighting and photographing objects of various shapes and reflective properties are well understood, being derived from optical physics and long practiced by artists and photographers. But the particular optical properties of smartphone touchscreens and the marks left behind on them are less well understood; we are aware of no comprehensive study or body of work that catalogs the conditions under which real-world smudges will or will not render well in photographs of such devices.

A comprehensive review of photographic lighting theory and practice is beyond the scope of this paper; an excellent tutorial can be found, for example, in [7]. What follows is a brief overview of the basic principles that underlie our experiments. In particular, we are concerned with several variables: the reflective properties of the screen and the smudge; the quality and location of the light sources; and finally, the location of the camera with respect to the screen.

Object surfaces react (or do not react) to light by either *reflecting* it or *diffusing* it. Reflective surfaces (such as mirrors) bounce light only at the complementary angle from which it arrived; an observer (or camera) sees reflected light only if it is positioned at the opposite angle. Diffuse surfaces, on the other hand, disperse light in all directions regardless of the angle at which it arrives; an observer will see diffused light at any position within a line of site to the object. The surfaces of most objects lie somewhere on a spectrum between being completely reflective and completely diffuse.

Lighting sources vary in the way they render an object's texture, depending on both the size and the angle of the light. The *angle* of the light with respect to the subject determines which surfaces of the object are highlighted and which fall in shadow. The *size* of the light with respect to the subject determines the range of angles that keep reflective surfaces in highlight and how shadows are defined. Small, point-size lights are also called *hard* lights; they render well-defined, crisp shadows. Larger light sources are said to be *soft*; they render shadows as gradients. Finally, the angle of the camera with respect to the subject surface determines the tonal balance between reflective and diffuse surfaces.

These standard principles are well understood. What is not well understood, however, is the reflective and diffuse properties of the screens used on smartphone devices or of the effects of finger smudges on these objects. We conducted experiments that varied the angle and size of lighting sources, and the camera angle, to determine the condition under which latent smudge patterns do and do not render photographically.

4.2 Photographic Setup

Our principle setup is presented in Fig. 3. We use a single light source (either soft, hard lighting, or omnidirectional lighting via a lighting tent) oriented vertically or horizontally. A vertical angle increments in plane with the camera, while a horizontal angle increments in a perpendicular plane to the camera. All angles are measured with respect to the smartphone.

Vertical angles were evaluated in 15 degree increments, inclusively between 15 and165 degrees. Degrees measures are complementary for vertical and lens angles. For example, a lens angle of 15 degrees and a vertical angle of 15 degrees are exactly complementary such that light reflects off the touch screen into the camera like a mirror. Horizontal angles were evaluated inclusively between 15 and 90 degrees as their complements produce identical effects. Similarly, we only consider camera angles between 15 and 90 degrees, inclusively; *e.g.*, a vertical and lens angle both at 105 degrees is equivalent to a vertical and lens angle both at 15 degrees with just the light and camera switch. Additionally, when the lens angle is at 90 degrees, only vertical lighting angles of 15 to 90 degrees need consideration[6]. Finally, for omnidirectional light only the lens angles need to be iterated as

[6] We do not consider 180 or 0 degree angles, which cannot provide lighting or exposure of the smudges.

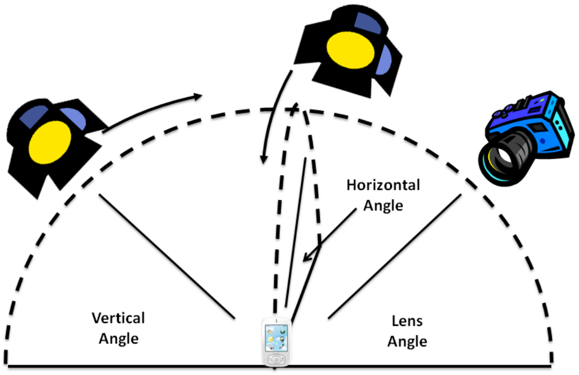

Figure 3: Principle Photographic Setup: The lighting and camera conditions at various vertical lighting angles (in plane with camera), horizontal lighting angles (in perpendicular plane with camera), and lens angles with respect to the smartphone.

light is dispersed such that it seems it is originating from all possible angles.

In total, there are 188 possible setups. For the base lighting condition, hard or soft, there are 11 vertical and 6 horizontal angles for 5 possible lens angles, not including the 90 degrees lens angle which only has 6 possible setups. With the addition of 6 lens angles for omnidirectional lighting, that leaves $188 = 2(5 \times 17 + 6) + 6$ setups, but there is still overlap. A 90 degree angle vertically and horizontally are equivalent, resulting in 178 unique setups.

4.3 Equipment Settings

We used relatively high end, precision cameras, lenses, lighting, and mounting equipment in our experiments to facilitate repeatability in our measurements. However, under real-world conditions, similar results could be obtained with much less elaborate (or expensive) equipment and in far less controlled environments.

All photographs were captured using a 24 megapixel Nikon D3x camera (at ISO 100 with 16 bit raw capture) with a tilting lens (to allow good focus across the entire touch screen plane). The camera was mounted on an Arca-Swiss C-1 precision geared tripod head. The large ("soft") light source was a 3 foot Kino-Flo fluorescent light panel; the small ("hard") light was a standard cinema "pepper" spotlight. For single light experiments, the directional light was at least 6 stops (64 times) brighter than ambient and reflected light sources. For omnidirectional lighting, we used a Wescott light tent, with light adjusted such that there was less than a 1 stop (2x) difference between the brightest and the dimmest light coming from any direction. All images were exposed based on an incident light reading taken at the screen surface.

4.4 Pattern Selection and Classification

In all experiments, we consider a single pattern for consistency, presented in Fig. 2. We choose this particular pattern because it encompasses all orientation and nearly all directions, with the exception of a vertical streak upwards. The direction and orientation of the pattern plays an important role in partial information collection. In certain cases, one direction or orientation is lost (see Sec. 6).

When determining the effectiveness of pattern identification from smudges, we use a simple classification scheme. First, two independent ratings are assigned on a scale from 0 to 2, where 0 implies that no pattern information is retrievable and 2 implies the entire pattern can be identified. When partial information about the pattern can be observed, *i.e.*, there is clearly a pattern present but not all parts are identifiable, a score of 1 is applied. Next, the two independent ratings are combined; we consider a pattern to be fully identifiable if it received a rating of 4, *i.e.*, both classifiers indicated full pattern extraction[7].

We also wished to consider the full extent of an attacker, so we allow our classifiers to adjust the photo in anyway possible. We found that with a minimal amount of effort, just by scaling the contrast slighting, a large number of previously obscured smudges become clear. Additionally, all the image alterations performed are equivalent to varying exposure or contrast settings on the camera when the image was captured.

5 Experiments

In this section, we present our experiments to test the feasibility of a smudge attack via photography. We conducted three experiments: The first considers ideal scenarios, where the touch screen is clean, and investigated the angles of light and camera that produce the best latent images. The results of the first experiment inform the later ones, where we simulate application usage and smudge removal based on contact with clothing.

5.1 Experiment 1: Ideal Collection

The goal of this experiment was to determine the conditions by which an attacker can extract patterns, and the best conditions, under ideal settings, for this. We consider various lighting and camera angles as well as different styles of light.

Setup. In this experiment we exhaust all possible lighting and camera angles. We consider hard and soft lighting as well as completely disperse, omnidirectional lighting, using a total of 188 photographs in classification. We

[7] We note that this rating system can lead to bias because the same pattern is used in every photograph. Specifically, there may be projection bias; knowing that a smudge streak is present, the classifier projects it even though it may not necessarily be identifiable. We use two independent classifiers in an attempt to alleviate this bias and only consider full pattern retrieval if bother classifiers rate with value 4.

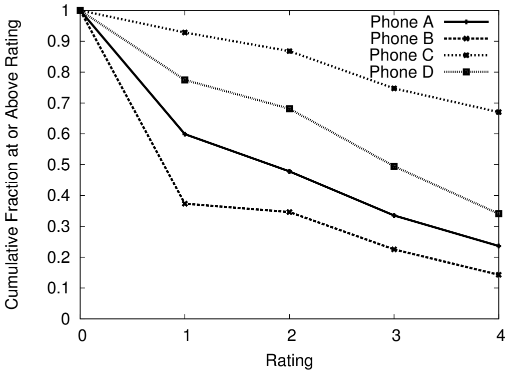

Figure 4: Cumulative Fraction Graph for Experiment 1: For each rating and phone, the cumulative fraction of photos scoring that rating, or higher.

	G1		Nexus 1	
App. Noise	over	under	over	under
dots	4	4	2.7	3.7
streaks	3	2	3	3
dots & steaks	3	1.6	4	3
face	4	2.3	4	2

Table 1: Results of Experiment 2: The average rating with application usage for patterns entered over and under the application noise.

experiment with four phones with different qualities of pattern entry, referred to by these letter identification:

- **Phone A:** HTC G1 phone with the pattern entered using "normal" touches
- **Phone B:** HTC G1 phone with the pattern entered using "light" touches
- **Phone C:** HTC G1 phone with the pattern entered after the phone has been held in contact with a face, as would happen after a phone call
- **Phone D:** HTC Nexus 1 phone with pattern entered using "normal" touches

Results. As described previously, each photograph is rated by the combination of two unique ratings on a scale from 0 to 2, which when combined provide a rating on a scale between 0 and 4. The key results of this classification are presented in Fig. 4 as a cumulative fraction graph.

The pattern that was most easily identifiable was Phone C, where the phone was first placed on a face prior to pattern entry. In roughly 96% of the photographic setups, a partial pattern was retrievable (*i.e.*, a rating of at least 1), and in 68% of the setups, the complete pattern was retrieved (*i.e.*, a rating of 4).

In contrast to the other tested phones, Phone C was dirty prior to password entry as broad smudging occurred due to contact with the facial skin. Entering the pattern on top of this broad smudge greatly contrasted with the pattern entry smudges (see Fig. A5). We explore this phenomenon further in Experiment 2. It is important to note that entering a pattern after a phone call is likely common because most conversations are longer than the phone lockout period. If a user wants access to other applications post hang-up, she will have to enter her pattern.

Phone B was the worst performing pattern entry. In this case, the pattern was entered using light touching, yet in over 30% of the setups, some partial information was retrievable. Moreover, in 14% of the photographs, the complete pattern is retrievable.

By far the best lens angle for retrieval was 60 degrees (followed closely by 45 degrees). In more than 80% of the lighting scenarios with a 60 degree lens, perfect or nearly perfect pattern retrieval was possible with a 60 degree camera angle. The worst retrieval was always when the vertical and lens angle were complimentary which transformed the touch screen surface into a mirror, effectively washing out the smudges (see Fig. A4 for one such example). Additionally, omnidirectional light (*i.e.*, using the light tent), had a similar effect. Omnidirectional light implies that there always exists a perfect reflection into the camera as light is emitted from all angles.

The most interesting observation made from the photographs is that in many of the setups, the *directionality* of the smudges can be easily discerned. That is, the order of the strokes can be learned, and consequently, the precise pattern can be determined. As an example see Fig. 5. At each direction change, a part of the previous stroke is overwritten by the current one, most regularly at contact points. On close inspection, the precise order of the contact points can be clearly determined and the pattern becomes trivially known.

5.2 Experiment 2: Simulated Usage

In this experiment, we were interested in the affect that user applications have on the capabilities of an attacker. Previously, we demonstrated that talking on the phone may increase the contrast between a pattern smudge and the background; we further elaborate on that point here. Additionally, we investigate the affect of application usage as it may occur prior to or post pattern entry.

Setup. The setup of this experiment was informed by the results of the previous. We photographed the phones

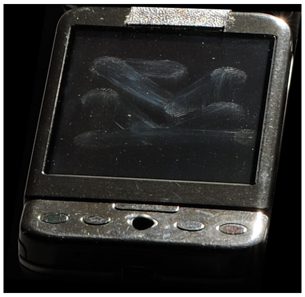

Figure 5: Phone A, from Experiment 1, where the pattern is entered with normal touches. Notice that the directionality of the pattern can be determined at ever direction change.

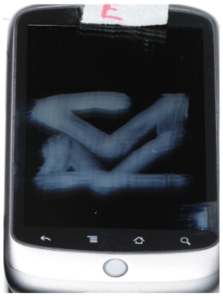

Figure 6: Phone from Experiment 3, where the phone was wiped, placed (and replaced) in a pocket. Unlike Phone A from Fig. 5, some directionality is lost in the upper left portion of the pattern.

at a 45 degree lens angle and at three of the best vertical angles: 15, 75, and 90 degrees[8].

We based our usage simulation on the telephone application; an application installed on all Android smartphones. Although the phone application is not representative of all application usage, it has some important characteristics present in nearly all applications. If a user were to enter a phone number it would require a sequence of presses, or *smudge dots*, on the screen. A user could also scroll her contact list, causing up and down *smudge streaking*, or left and right, depending on the phones current orientation. Additionally, there may be combinations of these.

For each phone in the experiment – two G1 phones and two Nexus 1 phones – we consider 3 usage scenarios; (1) A grid of smudge dotting; (2) a hashing of up and down and left right smudge streaks; and (3), a combination of smudge dots and streaks in a grid and hash, respectively. We also consider if the pattern was entered prior to application usage (*i.e.*, the pattern is first entered, and then the steaks and/or dots are applied) or post application usage (*i.e.*, first steaks and/or dots are applied followed by the pattern). Finally, as a follow up to Experiment 1, we consider the effect of placing the touch screen surface to the face, before and after pattern entry. In all pattern entries, we assume normal touching.

Results. As before, each photograph was classified by two individuals, and the combined results are considered. The results are summarized in Table 1. In general, entering the pattern over the usage smudges is more clearly retrieved, as expected. Dots also tend to have less of an effect then streaks, again as expected.

Interestingly, the over pattern entry for the combination of dots and streaks on the Nexus 1 scored perfect retrieval (see Fig. A1 for a sample image). Upon closer inspection, this is due to the intricacy of the pattern – the many hooks and turns required for such a long pattern – created great contrast with usage noise, and thus the pattern was more easily retrieved. Finally, as expected based on the results of Experiment 1, broad smudging on the face provided perfect retrieval for the over case, and even in the under case, partial information was retrieved.

5.3 Experiment 3: Removing Smudges

In this experiment we investigated the effects of smudge distortion caused by incidental contact with or wiping on clothing.

Setup. Using the same photographic setup as in Experiment 2, we photographed two clothing interference scenarios, both including placing and replacing the phone in a jeans pocket. In the first scenario, the user first intentionally wipes the phone, places it in her pocket, and removes it. In scenario two, the user places the phone in her pocket, sits down, stands up, and removes it.

Although this does not perfectly simulate the effects of clothing contact, it does provide some insight into the tenacity of a smudge on a touch screen. Clearly, a user can forcefully wipe down her phone until the smudge is no longer present, and such scenarios are uninteresting. Thus, we consider incidental distortion.

Results. Surprisingly, in all cases the smudge was classified as perfectly retrievable. Simple clothing contact does not play a large role in removing smudges. However, on closer inspection, information was being lost. The directionality of the smudge often could no longer be determine (See Fig. 6 for an example). Incidental wiping disturbed the subtle smudge overwrites that informed directionality. Even in such situations, an attacker has greatly reduced the likely pattern space to 2; the pattern in the forward and reverse direction.

[8] Although 60 degrees lens angle performed best overall, the setup required for 45 degrees was much simpler and had similarly good results at these vertical lighting angles.

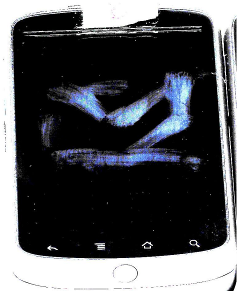

Figure 7: Phone from Experiment 1: One stroke of the pattern, |84|, is lost due to the camera or lighting angle. The contrast has been adjusted.

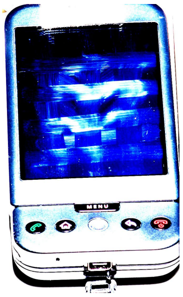

Figure 8: Phone from Experiment 2: With this usage condition (dot and streaks, under), the pattern is nearly all lost. The contrast has been adjusted.

5.4 Summary

Our photographic experiments suggest that a clean touchscreen surface is primarily, but not entirely, reflective, while a smudge is primarily, but not entirely, diffuse. We found that virtually any directional lighting source that is not positioned exactly at a complementary angle to the camera will render a recoverable image of the smudge. Very little photo adjustment is required to view the pattern, but images generally rendered best when the photo capture was overexposed by two to three f-stops (4 to 8 times "correct" exposure).

If the effect of the smudge is to make a chiefly reflective surface more diffuse, we would expect completely even omnidirectional light to result in very poor rendering of the image. And indeed, our experiments confirm this – even extensive contrast and color adjustment was generally unable to recover the smudge pattern from images captured under omnidirectional light under the light tent. Fortunately for the attacker, however, most "real world" lighting is primarily directional. The main problem for an attacker who wishes to surreptitiously capture a smudge pattern is not application noise or incidental clothing contact (as Experiment 2 and 3 showed) but rather ensuring that the angle of the camera with respect to the screen surface is not at an angle complementary to any strong light source.

6 Directions for Exploitation

We have demonstrated the ability of an attacker to capture information from smudges via photography. We now discuss how the information gained can be use to defeat the Android password pattern. As presented in Sec. 3, the size of the pattern space contains 389,112 distinct patterns. A significant number of those patterns can be eliminated as possible passwords by a smudge attacker. For example, perfect pattern retrieval with directionality is possible, reducing the possibilities to 1. Partial retrieval of patterns from smudges requires deeper analysis, towards which we present initial thoughts on exploiting captured smudges. Smudge data can be combined with statistical data on human behavior such as pattern usage distributions for large sets of users to produce likely sets of patterns for a particular smudged phone.

6.1 Using Partial Information

Using the photographs taken during our experiments, we investigated what was lost in partial retrieval scenarios. Two cases emerged: First, a lack of finger pressure and/or obscuration of regions of the photograph led to information loss. For example, in Fig. 7, the diagonal for connection |48| cannot be retrieved. This partial retrieval is still extremely encouraging for an attacker, who has learned a good deal about which patterns are likely, e.g., it could be each isolated part uniquely, the two parts connected, etc.

Another case emerges when a significant amount of usage noise obscures the smudge present; e.g., Fig. 8 is a photo from Experiment 2 with dots and streaks over the pattern entry. An attacker may guess that two sets of "V" style diagonals are present, but in general the entire pattern is not observable. Moreover, using this information is not likely to reduce the pattern space below the threshold of 20 guesses.

However, an attacker may have access to many images of the same pattern captured at different points in time. By combining this information, it may be possible for an attacker to recreate the complete pattern by data fusion [6]. As an example, consider an attacker combining the knowledge gained from the Fig. 7 and Fig. 8; if it was known that the same pattern was entered, the bottom "V" shape in Fig. 8 is enough information to finish the pattern in Fig. 7.

6.2 Human Factors

Human behavior in password selection has been well studied [11, 15], and even in the graphical password space, password attack dictionaries based on human mnemonics or pattern symmetry have been proposed [17, 18]. Similar behaviors seem likely to emerge in the Android password space, greatly assisting an attacker.

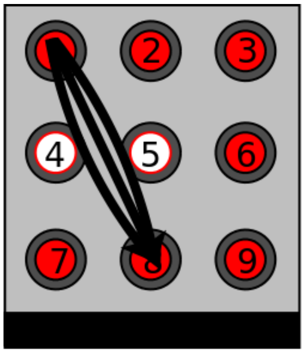

Figure 9: A 30 degree pattern stroke that is difficult to enter when points 4 and/or 5 are previously unselected.

We conjecture that the ease of pattern entry is an important factor for consideration because a user must enter her password on every phone lock; generally, a 30 second timeout. If the password is difficult to enter consistently, then it would be less usable and therefor less likely the user's chosen pattern. For example, the contact point stroke in Fig. 9 contains a 30 degree strokes which is prone to error when the intermittent contact points are not previously touched (*e.g.*, point 4 and 5). When considering this additional restriction, the password space can be reduced by over 50% to 158,410 likely patterns.

Another usability factor is pattern length: The longer the pattern, the longer the amount of time it takes to enter it. A frequent smartphone user may avoid excessively long patterns. In the same vein, a user may avoid frequent direction changes, as each requires, again, more time to enter. Other human factors may play a role in pattern selection and investigating them in the context of smudge attacks is an area of future research.

7 Related Work

Previous work on this subject is limited. Perhaps the closest related work was performed by Laxton *et al.* regarding copying physical keys based on photographic analysis [12], a so called *teleduplication attack*. An attacker may reproduce the key by measuring the position and relative depth of the key cuts to extract the bitting code. Once known, the key may be duplicated without the attacker ever having possessed it.

Smudge and teleduplication attacks are similar in a number of ways. First, both take advantage of visual information, whether password smudges or key bittings. Many of the same basic principles of teleduplication attacks, such as the photographic capturing methods, are relevant to our work. Both attacks are executed in physical space and can be done from afar. Finally, the usefulness of information gained requires next steps. In the case of a teleduplication attack, duplicating the key is only useful if the door it opens is known. In the same way, learning the pattern is only useful if the touch screen device were to come into the attacker's possession.

There has been a fair amount of interesting research performed on graphical passwords [2, 16]. Specifically, it should also be noted that there are several other proposed graphical password schemes [2, 4, 10]. We believe that several of these authentication procedures, if performed on a touch screen, may be susceptible to smudge attacks.

If smudge attacks were to be automated, previous work in the area automated image recognition, *e.g.* facial recognition techniques [5, 9] or optical character recognition [8, 13, 14], would be applicable. Such automated techniques are especially dangerous if an attacker possessed many successive images (*e.g.*, via video surveillance).

8 Conclusion

In this paper we explored *smudge attacks* using residual oils on touch screen devices. We investigated the feasibility of capturing such smudges, focusing on its effect on the password pattern of Android smartphones. Using photographs taken under a variety of lighting and camera positions, we showed that in many situations full or partial pattern recovery is possible, even with smudge "noise" from simulated application usage or distortion caused by incidental clothing contact. We have also outlined how an attacker could use the information gained from a smudge attack to improve the likelihood of guessing a user's patterns.

Next steps in our investigation include a deeper analysis of the advantages derived from analysis of smudges and an end-to-end smudge attack experiment on a small (voluntary) population of Android users who employ the password pattern. Additionally, we would like to perform a broad study of human factors in pattern selection as they relate to the Android pattern.

We believe smudge attacks based on reflective properties of oily residues are but one possible attack vector on touch screens. In future work, we intend to investigate other devices that may be susceptible, and varied smudge attack styles, such as heat trails [19] caused by the heat transfer of a finger touching a screen.

The practice of entering sensitive information via touch screens needs careful analysis in light of our results. The Android password pattern, in particular, should be strengthened.

Acknowledgments

We would like to thank the anonymous reviewers for their useful insight and thoughtful remarks, and Samual Panzer for his assistance during the experiments. Finally, this work was supported by *Biomedical Informa-*

tion Red Team (*BIRT*) NFS Award CNS-0716552 and *Security Services in Open Telecommunication Networks* NFS grant CNS-0905434.

References

[1] Android 2.2 platform highlights. http://developer.android.com/sdk/android-2.2-highlights.html.

[2] D. Davis, F. Monrose, and M. K. Reiter. On user choice in graphical password schemes. In *USENIX Sec'04*, 2004.

[3] A. M. DeAlvare. A framework for password selection. In *UNIX Security Workshop II*, 1998.

[4] H. Gao, X. Guo, X. Chen, L. Wang, and X. Liu. Yagp: Yet another graphical password strategy. *Computer Security Applications Conference, Annual*, 0:121–129, 2008.

[5] S. Gutta, J. R. Huang, H. Wechsler, and B. Takacs. Automated face recognition. volume 2938, pages 20–30. SPIE, 1997.

[6] D. L. Hall and J. Llinas. An introduction to multisensory data fusion. *Proc. IEEE*, 85(1), January 1997.

[7] F. Hunter and P. Fuqua. *Light: Science and Magic: An Introduction to Photographic Lighting*. Focal Press, 1997.

[8] S. Impedovo, L. Ottaviano, and S. Occhinegro. Optical character recognition – a survey. *International Journal of Pattern Recognition and Artificial Intelligence (IJPRAI)*, 5(1-2):1–24, 1991.

[9] R. Jenkins and A. Burton. 100% accuracy in automatic face recognition. *Science*, 319(5862):435, January 2008.

[10] I. Jermyn, A. Mayer, F. Monrose, M. K. Reiter, and A. D. Rubin. The design and analysis of graphical passwords. In *USENIX Sec'99*, pages 1–1, Berkeley, CA, USA, 1999. USENIX Association.

[11] D. V. Klein. Foiling the cracker: A survey of, and improvements to, password security. In *USENIX Sec'90*, 1990.

[12] B. Laxton, K. Wang, and S. Savage. Reconsidering physical key secrecy: Teleduplication via optical decoding. In *CCS*, October 2008.

[13] J. Mantas. An overview of character recognition methodologies. *Pattern Recognition*, 19(6):425–430, 1986.

[14] S. Mori, H. Nishida, and H. Yamada. *Optical Character Recognition*. John Wiley & Sons, Inc., New York, NY, USA, 1999.

[15] R. Morris and K. Thompson. Password security: a case history. *Communnincations of the ACM*, 22(11):594–597, 1979.

[16] K. Renaud and A. D. Angeli. Visual passwords: Cure-all or snake-oil. *Communications of the ACM*, 52(12):135–140, December 2009.

[17] J. Thorpe and P. van Oorschot. Graphical dictionaries and the memorable sapce of graphical passwords. In *USENIX Sec'04*, August 2004.

[18] J. Thorpe and P. C. van Oorschot. Human-seeded attacks and exploiting hot-spots in graphical passwords. In *USENIX Sec'07*, 2007.

[19] M. Zalewski. Cracking safes with thermal imaging, 2005. http://lcamtuf.coredump.cx/tsafe/.

A Appendix

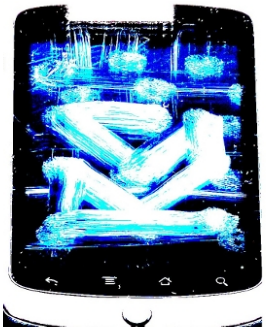

Figure A1: A phone from Experiment 2: The pattern contrasts greatly with the background noise; a grid of dots. The contrast on this image has been adjusted.

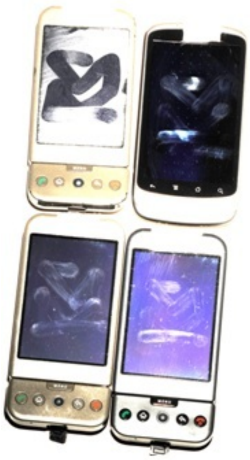

Figure A2: An image from Experiment 1: All four phones clearly displayed the pattern without the need to adjust contrast. Even the lightly touched Phone B has a visible pattern.

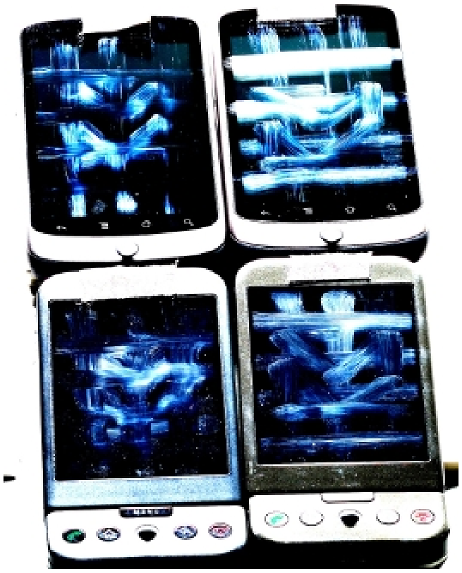

Figure A3: An image from Experiment 2: Even with background noise (over, on the left, and under, on the right of the pattern entry), either partial or complete pattern identification is possible as it contrast with such usage noise. The contrast on these images has been adjusted.

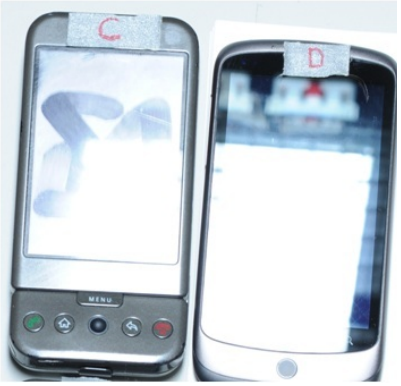

Figure A4: An image from Experiment 1: Complimentary lighting and lens angle causes significant glare, leading to unidentifiable patterns and information loss.

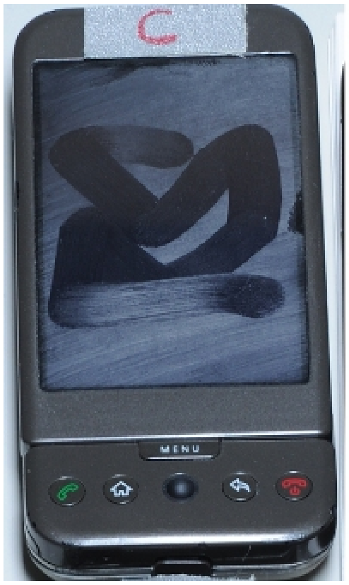

Figure A5: Phone C from Experiment 1: Without any image adjustment, the pattern is clear. Notice that the smudging, in effect, cleans the screen when compared to the broad smudging caused by facial contact. This contrast aids in pattern identification, as demonstrated in Experiment 2.

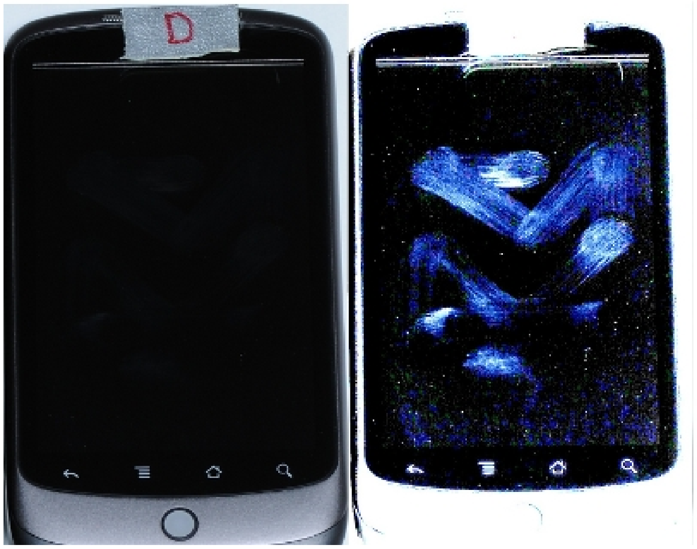

Figure A6: Phone D from Experiment 1, prior to and post contrast adjustment: In many situations, adjusting the levels of color or contrast can highlight a smudge previously obscured. The images on the left and right are identical.